arXiv:1703.08840v1 [cs.LG] 26 Mar 2017

# Inferring The Latent Structure of Human Decision-Making from Raw Visual Inputs

**Yunzhu Li** [1] **Jiaming Song** [2] **Stefano Ermon** [2]

## Abstract

The goal of imitation learning is to match example expert behavior, without access to a reinforcement signal. Expert demonstrations provided by humans, however, often show significant variability due to latent factors that are not explicitly modeled. We introduce an extension to the Generative Adversarial Imitation Learning method that can infer the latent structure of human decision-making in an unsupervised way. Our method can not only imitate complex behaviors, but also learn interpretable and meaningful representations. We demonstrate that the approach is applicable to high-dimensional environments including raw visual inputs. In the highway driving domain, we show that a model learned from demonstrations is able to both produce different styles of human-like driving behaviors and accurately anticipate human actions. Our method surpasses various baselines in terms of performance and functionality.

## 1. Introduction

A key limitation of reinforcement learning (RL) is that it involves the optimization of a predefined reward function or reinforcement signal (Levine & Koltun, 2013; Schulman et al., 2015a; Lillicrap et al., 2015; Schulman et al., 2015b; Silver et al., 2016; Tamar et al., 2016). Defining a reward function is straightforward in some cases, e.g., in games such as Go or chess. However, designing an appropriate reward function can be difficult in more complex and less well-specified environments, e.g., for autonomous driving where there is a need to balance safety and efficiency.

Imitation learning methods have the potential to close this gap by learning how to perform tasks directly from expert demonstrations, and has succeeded in a wide range of problems (Ziebart et al., 2008; Ratliff et al., 2009; Stadie et al., 2017). Among them, Generative Adversarial Imitation Learning (GAIL, Ho & Ermon (2016)) is a model-free imitation learning method that is highly effective and scales to relatively high dimensional models. The training process of GAIL can be thought of as building a generative model, a stochastic policy that reacts to a fixed simulation environment, to produce behavior that matches the expert demonstrations. To determine a reasonable distance metric, a discriminator is jointly trained to distinguish expert trajectories from ones produced by the policy.

In imitation learning, example demonstrations are typically provided by human experts. These demonstrations can show significant variability. For example, they might be collected from multiple experts, each employing a different policy. External latent factors of variation that are not explicitly captured by the simulation environment can also significantly affect the observed behavior. For example, expert driving demonstrations might be collected from users with different skills and habits. The goal of this paper is to develop an imitation learning framework that is able to automatically discover and disentangle the latent factors of variation underlying human decision-making. Analogous to the goal of uncovering style, shape, and color in generative modeling of images (Chen et al., 2016), we aim to automatically learn concepts such as driver aggressiveness from human demonstrations.

We propose a new method for learning a latent variable generative model of trajectories in a dynamic environment that not only accurately reproduce expert behavior, but also learns a latent space that is semantically meaningful. Our approach is an extension of GAIL, where the objective is augmented with a mutual information term between the latent variables and the observed state-action pairs. We demonstrate an application in autonomous driving, where we learn to imitate complex driving behaviors while learning semantically meaningful structure, without any supervision beyond the expert trajectories. Remarkably, our method performs directly on raw visual inputs, using raw pixels as the only source of perceptual information.

In particular, the contributions of this paper are threefold:

[1]Peking University, Beijing, China [2]Stanford University, California, USA. Correspondence to: Yunzhu Li <leo.liyunzhu@pku.edu.cn>, Stefano Ermon <ermon@cs.stanford.edu>.

1. We introduce a component which maximizes the mutual information between latent structure and trajectories, similar to InfoGAN (Chen et al., 2016) , resulting in a policy where low-level actions can be controlled through high-level latent variables.

2. We extend GAIL to use raw pixels as input and produce human-like behaviors in complex high-dimensional environments.

3. We demonstrate an application to autonomous highway driving using the TORCS driving simulator (Wymann et al., 2000). We first demonstrate that the learned policy is able to navigate traffic without collisions. Then, we show that by specifying high-level latent variables, our model learns to reproduce different styles of human-like driving behavior.

## 2. Background

### 2.1. Preliminaries

We use the tuple $(\mathcal{S}, \mathcal{A}, P, r, \rho_0, \gamma)$ to define an infinite-horizon, discounted Markov decision process (MDP), where $\mathcal{S}$ represents the state space, $\mathcal{A}$ represents the action space, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \mathbb{R}$ denotes the transition probability distribution, $r : \mathcal{S} \to \mathbb{R}$ denotes the reward function, $\rho_0 : \mathcal{S} \to \mathbb{R}$ is the distribution of the initial state $s_0$, and $\gamma \in (0, 1)$ is the discount factor. Let $\pi$ denote a stochastic policy $\pi : \mathcal{S} \times \mathcal{A} \to [0, 1]$, and $\pi_E$ denote the expert policy to which we only have access to demonstrations. The expert demonstrations $\tau_E$ are a set of trajectories generated using policy $\pi_E$, each of which consists of a sequence of state-action pairs.

### 2.2. Imitation learning

The goal of imitation learning is to learn how to perform a task directly from expert demonstrations, without any access to the reinforcement signal $r$. Typically, there are two approaches to imitation learning: 1) behavior cloning (BC), which learns a policy through supervised learning over the state-action pairs from the expert trajectories; and 2) apprenticeship learning (AL), which assumes the expert policy is optimal under some unknown reward and learns a policy by recovering the reward and solving the corresponding planning problem. BC tends to have poor generalization properties due to compounding errors and covariate shift (Ross & Bagnell, 2010; Ross et al., 2011). AL, on the other hand, has the advantage of learning a reward function that can be used to score entire trajectories (Abbeel & Ng, 2004; Syed et al., 2008; Ho et al., 2016), but is typically expensive to run because it requires solving a reinforcement learning (RL) problem inside a learning loop.

### 2.3. Generative Adversarial Imitation Learning

Recent work on AL has adopted a different approach by learning a policy without learning the reward function. In particular, Generative Adversarial Imitation Learning (GAIL, Ho & Ermon (2016) ) is a recent AL method inspired by Generative Adversarial Networks (GAN, Goodfellow et al. (2014)). In the GAIL framework, the agent imitates the behavior of an expert policy $\pi_E$ by matching the generated state-action distribution with the expert distribution, where the optimum is achieved when these two distributions match perfectly. However, measuring the similarity between high-dimensional distributions is complicated, so GAIL introduces a neural network to approximately minimize the Jensen-Shannon divergence. Intuitively, the neural network is a discriminator that tries to differentiate the two distributions. The formal GAIL objective is denoted as $\min_\theta \max_\omega V(\theta, \omega)$, where $V(\theta, \omega)$ is

$$\mathbb{E}_{\pi_\theta}[\log D_\omega(s, a)] + \mathbb{E}_{\pi_E}[\log(1 - D_\omega(s, a))] - \lambda H(\pi_\theta) \quad (1)$$

and $\pi_\theta$ (usually a neural network parameterized by $\theta$) is the policy that we wish to imitate $\pi_E$ with, $D_\omega$ is a discriminator network which tries to distinguish state-action pairs from the trajectories of $\pi_\theta$ and $\pi_E$, $\mathbb{E}_\pi[f(s, a)]$ denotes the expectation of $f$ over the state-action pairs generated by $\pi$, and $H(\pi_\theta) \triangleq \mathbb{E}_{\pi_\theta}[-\log \pi_\theta(a|s)]$ is the $\gamma$-discounted causal entropy of the policy $\pi_\theta$ (Bloem & Bambos, 2014). We assume that our policies are Gaussian distributions with fixed standard deviations, thus $H(\pi_\theta)$ is constant in our settings.

Instead of directly learning a reward function, GAIL relies on the discriminator to guide $\pi_\theta$ into imitating the expert policy. Optimization over the GAIL objective is performed by alternating between an Adam (Kingma & Ba, 2014) gradient step on $\omega$ to increase $V(\theta, \omega)$ with respect to $D$, and a Trust Region Policy Optimization (TRPO, Schulman et al. (2015a)) step on $\theta$ to decrease $V(\theta, \omega)$ with respect to $\pi$. GAIL is model-free: it requires a simulator to provide rollouts, but it does not need to construct a model for the environment.

## 3. Visual InfoGAIL

Expert demonstrations are typically collected from human experts. The resulting trajectories can show significant variability among different individuals due to internal latent factors of variation, such as social values, character, or mood. In this section, we propose an approach that 1) can discover and disentangle salient latent factors underlying human decision-making without supervision, 2) learns policies that produce trajectories which correspond to these latent factors, and 3) uses visual inputs as the only external perceptual information.

Formally, we assume that the expert policy is a mixture of experts, and we define the generative process of the expert trajectory $\tau_E$ as: $s_0 \sim \rho_0$, $c \sim p(c)$, $\pi \sim p(\pi|c)$, $a_t \sim \pi(a_t|s_t)$, $s_{t+1} \sim P(s_{t+1}|a_t, s_t)$ where $c$ is a latent variable that selects a specific policy $\pi$ from the mixture of expert policies through $p(\pi|c)$, and $p(c)$ is the prior distribution of $c$. Similar to the settings in GAIL, we consider the apprenticeship learning problem as a dual of an occupancy measure matching problem, and treat the trajectory $\tau$ as a set of state-action pairs. The objective is to recover $\pi(a|s, c)$ from the state-action pairs given the prior $p(c)$.

### 3.1. Interpretable Imitation Learning

We model our parametrized policy as $\pi_\theta(a|s, c)$. To draw a connection between latent codes $c$ and the behavior of $\pi_\theta$, we utilize the information-theoretic regularization that there should be high mutual information between $c$ and the state-action pairs in the generated trajectory. This concept is introduced by InfoGAN (Chen et al., 2016), where latent codes are utilized to discover the salient structured semantic features of the data distribution and guide the generating process. In particular, the regularization seeks to maximize the mutual information between latent codes and state-action pairs, denoted as $I(c; s, a)$, which is hard to maximize directly as it requires access to the posterior. Hence we introduce a variational lower bound, $L_I(\pi_\theta, Q)$, of the mutual information $I(c; s, a)$[1]:

$$\begin{aligned} L_I(\pi_\theta, Q) &= \mathbb{E}_{c\sim p(c), a\sim\pi_\theta(\cdot|s,c)}[\log Q(c|s, a)] + H(c) \\ &\leq I(c; s, a) \end{aligned} \tag{2}$$

where $Q(c|s, a)$ is an approximation of $P(c|s, a)$ parameterized with weights $\psi$. The objective under this regularization, which we call Information Maximizing Generative Adversarial Imitation Learning (InfoGAIL), then becomes:

$$\min_{\theta,\psi} \max_{\omega} V(\theta, \omega) - \lambda_1 L_I(\pi_\theta, Q_\psi) \tag{3}$$

where $\lambda_1 > 0$ is the hyperparameter for information maximization regularization.

By introducing the latent code, InfoGAIL is able to identify the salient factors in the expert trajectories through mutual information maximization, and imitate the corresponding expert policy through generative adversarial training. This allows us to disentangle trajectories that may arise from a mixture of experts, such as different individuals performing the same task. We note that $L_I(\pi_\theta, Q_\psi)$ can be optimized through stochastic gradient methods, with $\pi_\theta$ updated by TRPO, and $Q_\psi$ updated by Adam.

[1] Chen et al. (2016) presents a proof for the lower bound.

### 3.2. Utilizing Raw Visual Inputs via Transfer Learning

In many real world applications, the state $s$ is related to visual inputs, such as an image or a sequence of images; indeed, visual data is often inexpensive to obtain, highly informative, and heavily relied upon by people when performing tasks. Although our approach is general, we will focus on scenarios where the states $s$ are represented by images. Despite recent successes in visual perception, this is a very challenging scenario as raw visual inputs are typically very high-dimensional. As a result, learning a policy mapping high-dimensional raw visual inputs to actions is particularly difficult. Intuitively, the policy will have to simultaneously learn how to identify meaningful visual features, and how to leverage them to achieve desired behavior.

Convolutional neural networks (CNNs) have led to dramatic improvements across many computer vision tasks (Krizhevsky et al., 2012). Unfortunately, they require very large amounts of training data; therefore, using the trajectories to directly train the policy and image recognition will be expensive. Hence, methods to mitigate the high sample complexity problem are crucial to the success of InfoGAIL on raw visual inputs.

Here, we take a transfer learning approach. Features extracted using a CNN pre-trained on ImageNet contain high-level information about the input images, which can be adapted to new vision tasks via transfer learning (Yosinski et al., 2014). However, it is not yet clear whether these relatively high-level features can be directly applied to tasks where perception and action are tightly interconnected.

We show that it is indeed possible. We perform transfer learning by exploiting features from a pre-trained neural network that effectively convert raw images into relatively high-level information (Sharif Razavian et al., 2014). In particular, we use a Deep Residual Network (He et al., 2016) pre-trained on the ImageNet classification task (Russakovsky et al., 2015) to obtain the visual features used as inputs for the policy network.

## 4. Improved Optimization

While GAIL is successful in tasks with low-dimensional inputs (in Ho & Ermon (2016), the largest observation has 376 continuous variables), few have explored tasks where the input dimension is very high (such as $110 \times 200 \times 3$ pixels as in our experiments), even with pre-trained features from Residual Networks. In order to effectively learn a policy that relies solely on visual input, we make the following improvements over the original GAIL framework.

### 4.1. Reward Augmentation

In complex and less well-specified environments, imitation learning methods have the potential to perform better than reinforcement learning methods since they do not require manual specification of an appropriate reward function. Assuming the expert is following an optimal policy under some unknown reward, and that the expert is optimal, the reward function learned from the expert trajectories should match the desired reinforcement signal. However, if the expert is performing sub-optimally, then reward functions recovered from the expert will be imperfect. Therefore, any policy trained under these learned rewards will be also suboptimal; in other words, the imitation learning agent's potential is bounded by the capabilities of the expert.

In many cases, it is very difficult to fully specify a suitable reward function for a given task, yet, it is relatively straightforward to come up with constraints that we would like to enforce over the policy. For example, even if we cannot design an appropriate reward function capturing the subtleties of autonomous driving, we know that an autonomous vehicle should avoid colliding with other objects.

This motivates the introduction of *reward augmentation*, a general framework to incorporate prior knowledge in imitation learning by providing additional incentives to the agent without interfering with the imitation learning process. We achieve this by specifying a surrogate state-based reward $\eta(\pi_\theta) = \mathbb{E}_{s\sim\pi_\theta}[r(s)]$ that reflects our biases over the desired agent's behavior:

$$\min_\theta \max_\omega V(\theta,\omega) - \lambda_0 \eta(\pi_\theta) \tag{4}$$

where $\lambda_0 > 0$ is a hyper-parameter. This approach can be seen as a hybrid between imitation and reinforcement learning, where part of the reinforcement signal for the policy optimization is coming from the surrogate reward and part from the discriminator, i.e., from mimicking the expert. The surrogate reward can also be thought of as side information provided to the generator. For example, in our autonomous driving experiment below we show that by providing the agent with a penalty if it collides with other cars, we are able to significantly reduce the collision rate of the policy.

Computationally, this does not interfere with discriminator training, and for the policy gradient algorithm the surrogate reward is simply an additional gradient term for the TRPO update (Schulman et al., 2015a).

### 4.2. Wasserstein GAN

Wasserstein GAN (WGAN, Arjovsky et al. (2017)) is a recently proposed framework for generative adversarial training. Unlike its traditional GAN counterpart, the discriminator network in WGAN solves a regression problem instead of a classification problem by assigning scores to its inputs, and tries to maximize the score of real data while minimizing that of generated data. We extend the use of WGAN to the generative adversarial imitation learning framework, defining a new objective $W(\theta,\omega)$:

$$\mathbb{E}_{\pi_\theta}[D_\omega(s,a)] - \mathbb{E}_{\pi_E}[D_\omega(s,a)] \tag{5}$$

Extending the analysis of (Arjovsky et al., 2017), it can be shown that if $D$ is $K$-Lipschitz, the objective approximately minimizes the Earth-Mover (EM) distance between the distribution of trajectories from $\pi_\theta$ and $\pi_E$. This objective function suffers less from the vanishing gradient and mode collapse problems compared to traditional GANs. This is especially important in our setting, where we want to model complex distributions over trajectories that can potentially have a large number of modes. In Arjovsky et al. (2017), the $K$-Lipschitz property is satisfied by clipping weights in $D_\omega$ to between $[-0.01, 0.01]$, and using momentum-free optimization methods such as RMSProp (Tieleman & Hinton, 2012).

### 4.3. Variance Reduction

Policy gradients methods are notorious for suffering from high-variance gradients, since it is computationally expensive to obtain enough rollouts from the simulator to match the policy distribution. Therefore, we apply several variance-reduction techniques, such as replay buffer (Schaul et al., 2015) and baseline methods (Williams, 1992).

### 4.4. Algorithm

Apart from the baseline, we have three networks to update in the InfoGAIL framework: the discriminator network $D_\omega(s,a)$, the policy network $\pi_\theta(a|s,c)$, and the posterior estimator network $Q_\psi(c|s,a)$. We update $D_\omega$ using RMSprop (as suggested in the original WGAN paper), and update $Q_\psi$ and $\pi_\theta$ using Adam and TRPO respectively. The training procedure is presented in Algorithm 1. To speed up training, we initialize our policy from a behavior cloning policy, as in Ho & Ermon (2016).

The discriminator network $D_\omega$ and the posterior approximation network $Q_\psi$ are treated as distinct networks, as opposed to the InfoGAN approach where they share the same network parameters until the final output layer. This is because $D_\omega$ requires weight clipping and momentum-free optimization methods, which is required by the current WGAN training framework. These changes would interfere with the training of an expressive $Q_\psi$ if $D_\omega$ and $Q_\psi$ share the same network parameters.

**Algorithm 1** InfoGAIL

**Input:** Expert trajectories $\tau_E \sim \pi_E$; initial policy, discriminator and posterior parameters $\theta_0, \omega_0, \psi_0$; replay buffer $B = \varnothing$;
**Output:** Learned policy $\pi_\theta$
**for** $i = 0, 1, 2, ...$ **do**
Sample a batch of latent codes: $c_i \sim P(c)$
Sample trajectories: $\tau_i \sim \pi_{\theta_i}(c_i)$, with the latent code fixed during each rollout.
Update the replay buffer: $B \leftarrow B \cup \tau_i$.
Sample $\chi_i \sim B$ and $\chi_E \sim \tau_E$ with same batch size.
Update $\omega_i$ by ascending with gradients

$$\Delta_\omega = \hat{\mathbb{E}}_{\chi_i}[\nabla_\omega D_\omega(s,a)] - \hat{\mathbb{E}}_{\chi_E}[\nabla_\omega(D_\omega(s,a))]$$

Clip the weights of $\omega_i$ to $[-0.01, 0.01]$.
Update $\psi_{i+1}$ by descending with gradients

$$\Delta_\psi = -\lambda_1 \hat{\mathbb{E}}_{\chi_i}[\nabla_\psi \log Q_\psi(c|s,a)]$$

Take a policy step from $\theta_i$ to $\theta_{i+1}$, using the TRPO update rule with the following objective (without reward augmentation):

$$\hat{\mathbb{E}}_{\chi_i}[D_{\omega_{i+1}}(s,a)] - \lambda_1 L_I(\pi_{\theta_i}, Q_{\psi_{i+1}})$$

or (with reward augmentation):

$$\hat{\mathbb{E}}_{\chi_i}[D_{\omega_{i+1}}(s,a)] - \lambda_0 \eta(\pi_{\theta_i}) - \lambda_1 L_I(\pi_{\theta_i}, Q_{\psi_{i+1}})$$

**end for**

## 5. Experiments

We demonstrate the performance of our method by applying it to a complex autonomous driving from visual inputs domain. By conducting experiments on a car racing simulator, we show that our learned policy $\pi_\theta$ can 1) imitate human behavior using raw visual input using only a handful of expert demonstrations, 2) cluster human behaviors into different and semantically meaningful categories, and 3) reproduce different styles of human-like driving behaviors by setting the high-level latent variables.

### 5.1. Environment Setup

The Open Racing Car Simulator (TORCS, Wymann et al. (2000)) is a popular simulator environment for research in autonomous vehicles. We packaged it into a client-server framework with APIs similar to OpenAI Gym (Brockman et al., 2016). Our framework produces a realistic dashboard view and driving related information, and communicates with the policy (client) through TCP packets, so that the policy can be written in languages other than C++. In particular, we implemented our policy using the TensorFlow Python API (Abadi et al., 2016). This framework and the code for reproducing the experiments are available at `https://github.com/YunzhuLi/InfoGAIL`.

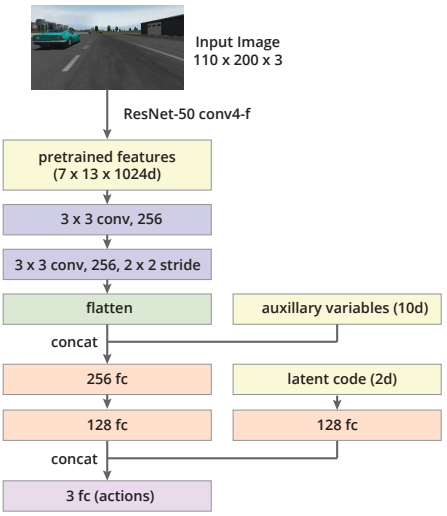


*Figure 1.* Network architecture for the policy/generator $\pi_\theta$. *conv* denotes a convolutional layer, and *fc* denotes a fully connected layer.

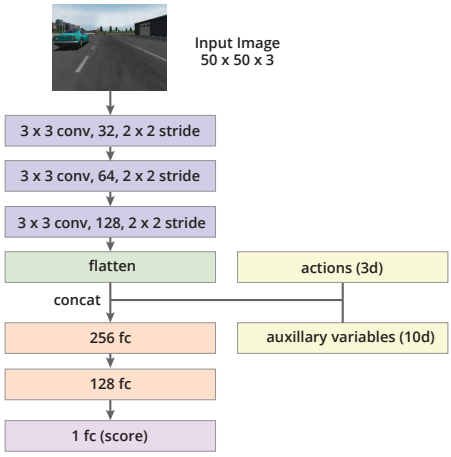
Input Image
50 x 50 x 3
3 x 3 conv, 32, 2 x 2 stride
3 x 3 conv, 64, 2 x 2 stride
3 x 3 conv, 128, 2 x 2 stride
flatten
actions (3d)
concat
256 fc
auxillary variables (10d)
128 fc
1 fc (score)

*Figure 2.* Network architecture for the discriminator $D_\omega$.

All of our experiments are conducted in the TORCS environment. The demonstrations are collected from human experts, by manually driving along the race track, and demonstrate typical behaviors like staying within lanes, avoiding collisions with other cars, and surpassing other cars. The policy accepts raw visual inputs as the only external inputs for the state, and produces a three-dimensional action that consists of *steering*, *acceleration*, and *braking*.

### 5.2. Network Structure

In addition, our policy requires certain auxiliary information as internal input to serve as a short-term memory. These auxiliary information can be accessed along with the

raw visual inputs. In our experiments, the auxiliary information for the policy at time $t$ consists of the following: 1) **velocity** at time $t$, which is a three dimensional vector; 2) **actions** at time $t-1$ and $t-2$, which are both three dimensional vectors; 3) **damage** of the car, which is a real value. The auxiliary input has 10 dimensions in total.

For the policy network, input visual features are passed through two convolutional layers, and then combined with the auxiliary information vector and (in the case of InfoGAIL) the latent code $c$. The exact architecture for $\pi_\theta$ is in Figure 1. Moreover, merging latent codes at the higher levels would have less effect over the actions, since the visual features have much larger dimensions. We parameterize the baseline as a network with the same architecture except for the final layer, which is just a one dimensional output indicates the expected accumulated future rewards.

The discriminator $D_\omega$ accepts three elements as input: a resized image with lower resolution, the auxiliary information, and the current action. The output is a score for the WGAN training objective, which is supposed to be higher for expert state-action pairs, and lower for generated ones. Details for the architecture of $D_\omega$ are shown in Figure 2. The posterior approximation network $Q_\psi$ adopts the same architecture as the discriminator except that the output is a softmax over the discrete latent variables, or factored Gaussian over continuous latent variables.

### 5.3. Inferring The Latent Structure of Human Decision-Making

In this experiment, we consider two subsets of human driving behaviors: ***turn***, where the agent takes a turn using either the inside lane or the outside lane; and ***pass***, where the agent passes another vehicle from either the left or the right. In both cases, the expert policy has two significant modes. Our goal is to have InfoGAIL capture the two modes from expert demonstrations.

We consider using a discrete latent code, which is a one-hot encoded vector with two possible states. For both settings, there are 80 expert trajectories in total, with 100 frames in each trajectory. The performance of a learned policy is quantified with two metrics: the *average distance* is determined by the distance traveled by the agent before a collision (and is bounded by the length of the simulation horizon), and *accuracy* is defined as the classification accuracy of the expert state-action pairs according to the latent code inferred with $Q_\psi$.

The average distance and sampled trajectories at different stages of training are shown in Figures 3 and 4 for ***turn*** and ***pass*** respectively. During the initial stages of training, the model does not distinguish the two modes and has a high chance of colliding, due to the limitations of behavior cloning (which we used to initialize the policy). As training progresses, trajectories provided by the learned policy begin to diverge. Towards the end of training, the two types of trajectories are clearly distinguishable, with only a few exceptions. In ***turn***, $[0, 1]$ corresponds to using the inside lane, while $[1, 0]$ corresponds to the outside lane. In ***pass***, the two kinds of latent codes are corresponding to passing from right and left separately. Meanwhile, the average distance of the rollouts steadily increases with more training.

*Table 1.* **Accuracy of unsupervised and supervised methods for *pass*.** Unsupervised methods are tested over all the expert trajectories; supervised methods are trained over $70\%$ of the trajectories and tested over the remaining $30\%$.

| UNSUPERVISED | | SUPERVISED | |
|---|---|---|---|
| K-MEANS | 55.4% | SVM | 85.8% |
| PCA | 61.7% | CNN | 90.8% |
| **INFOGAIL (OURS)** | **81.9%** | | |

In Figure 5, we show the visual input for the policy at certain time steps, during which the agent observes itself passing a corner or surpassing another vehicle from the side determined by the latent code. This behavior is visually similar to the corresponding expert behavior.

Learning the two modes separately requires accurate inference of the latent code. To examine the accuracy of posterior inference, we select state-action pairs from the expert trajectories and obtain the corresponding latent code through $Q_\psi(c|s, a)$. See Table 1, although we did not explicitly provide labeling, our model is able to correctly distinguish over $81\%$ of the state-action pairs in ***pass*** (and almost all the pairs in ***turn***, which can be seen in Figure 3).

For comparison, we also visualize the trajectories of ***pass*** for the original GAIL objective in Figure 4, where there is no maximum mutual information regularization. GAIL learns the expert trajectories as a whole, and cannot distinguish the two modes in the expert policy.

Interestingly, instead of learning two separate trajectories, GAIL tries to fit the left trajectory by swinging the car suddenly to the left *after* it has surpassed the other car from the right. We believe this reflects a limitation in the discriminators. Since $D_\omega(s, a)$ only requires state-action pairs as input, the policy is only required to match most of the state-action pairs; matching each rollout with a particular expert trajectory is not necessary. InfoGAIL with discrete latent codes can alleviate this problem by forcing the model to learn separate trajectories.

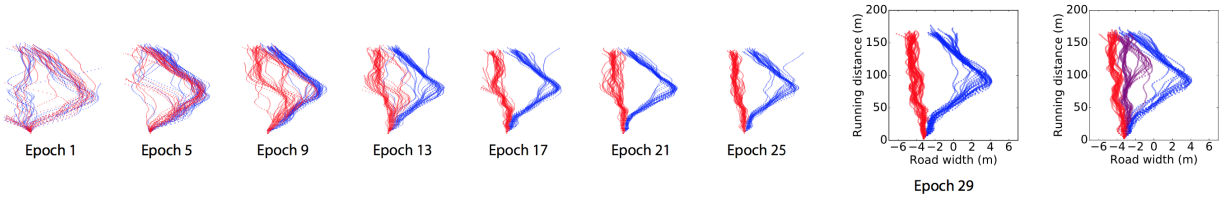


*Figure 3.* **Visualizing the training process of *turn*.** Here we show the trajectories of InfoGAIL at different stages of training. Blue and red indicates policies under different latent codes. The axes are the traveling distance and horizontal positions with respect to the road. Cutting a corner requires moving from one side to the other, and then moving back. Hence, blue indicates turning from the inner lane, and red indicates turning from the outer lane. The rightmost figure shows the trajectories under latent codes $[1, 0]$ (red), $[0, 1]$ (blue), and $[0.5, 0.5]$ (purple), which suggests that, to some extent, our method is able to generalize to cases where no trajectories are provided. Videos of these three kinds of trajectories are provided in the supplementary materials.

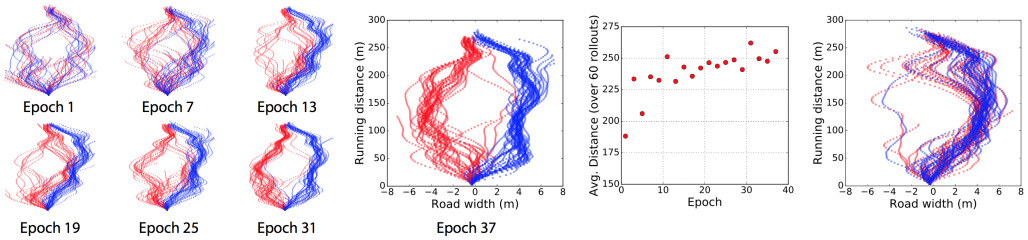


*Figure 4.* **Experimental results for *pass*. Left**: Trajectories of InfoGAIL at different stages of training. Blue and red indicates policies under different latent codes. **Middle**: Running distance denotes the absolute distance from the start position and is long since the other car is also moving quickly. Average distance over 60 rollouts of the InfoGAIL policy trained at different epochs. **Right**: Trajectories of ***pass*** produced by an agent trained on the original GAIL objective.

*Table 2.* **Average rollout distances of different policies.**

| METHOD | AVG. ROLLOUT DISTANCE |
|---|---|
| BEHAVIOR CLONING | 701.83 |
| INFOGAIL \ RB | 1031.13 |
| INFOGAIL \ RA | 1123.89 |
| INFOGAIL \ WGAN | 1177.72 |
| **INFOGAIL (OURS)** | **1226.68** |
| HUMAN | 1203.51 |

### 5.4. Ablation Experiments

We conduct a series of ablation experiments to demonstrate that our proposed techniques are indeed crucial for learning an effective policy. The experiments consider a long-term setting: our policy drives a car on the race track along with other cars, whereas the human expert provides trajectories by trying to drive as fast as possible without collision. Reward augmentation is performed by adding a reward that encourages the car to drive faster to the imitation learning objective. The performance of the policy is determined by the average distance. Therefore, a longer average rollout distance indicates a better policy.

In our ablation experiments, we remove parts of the improved optimization methods in Section 4. **InfoGAIL(Ours)** includes all the optimization techniques; **InfoGAIL\WGAN** switches the WGAN objective with the GAN objective; **InfoGAIL\RA** removes the reward augmentation term from the objective; **InfoGAIL\RB** removes the replay buffer and only samples from the most recent rollouts; **Behavior Cloning** is the behavior cloning method we use for initialization, and **Human** is the expert policy.

Table 2 shows the average rollout distances of different policies. Our method is able to outperform the expert with the help of reward augmentation; policies without reward augmentation or WGANs perform slightly worse than the expert; removing the replay buffer causes the performance to deteriorate significantly due to increased variance in gradient estimation.

## 6. Related work

There are two major paradigms for vision-based driving systems (Chen et al., 2015). *Mediated perception* is a two-step approach that first obtains scene information and then makes a driving decision (Aly, 2008; Lenz et al., 2011); *behavior reflex*, on the other hand, adopts a direct approach by mapping visual inputs to driving actions (Pomerleau, 1989; 1991). Many of the current autonomous

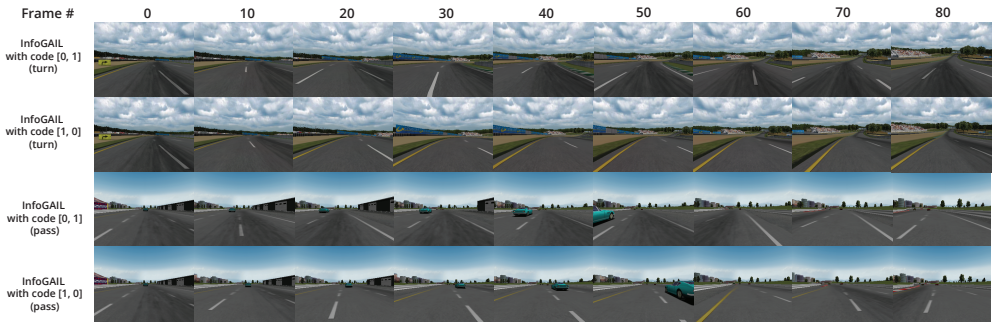


*Figure 5.* Visual inputs for the InfoGAIL policies in ***pass*** and ***turn*** experiments. The latent code controls the behavior of the policies. Videos are included in the supplementary material.

*Figure 6.* **Visual inputs for our policy in the long-term experiment.** During the rollout, our policy is able to pass several other cars. A video is included in the supplementary material.

driving methods rely on the two-step approach, which requires hand-crafting features such as the detection of lane markings and cars (Geiger et al., 2013; Chen et al., 2015). Our approach, on the other hand, attempts to learn these features through end-to-end training. While mediated perception approaches are currently more prevalent, we believe that end-to-end learning methods are more scalable and may lead to better performance in the long run.

Bojarski et al. (2016) introduce an end-to-end imitation learning framework that learns to drive entirely from visual information, and test their approach on real-world scenarios. However, their method uses behavior cloning by performing supervised learning over the state-action pairs, which is well-known to generalize poorly to more sophisticated tasks, such as changing lanes or passing vehicles. With the use of GAIL, our method can learn to perform these sophisticated operations easily.

Most imitation learning methods for end-to-end driving rely heavily on LIDAR-like inputs to obtain precise distance measurements (Ho et al., 2016; Kuefler et al., 2017). These inputs are not usually available to humans during driving. In particular, Kuefler et al. (2017) applies GAIL to the task of modeling human driving behavior on highways. Their policy is modeled using a recurrent neural network, which is supposed to maintain sufficient statistics of the past observations. In contrast, our policy requires only *raw visual information* as external input, which in practice is all the information humans need in order to drive.

Sermanet et al. (2016) have also introduced a pre-trained deep neural network to achieve better performance in imitation learning with relatively few demonstrations. Specifically, they introduce a pre-trained model to learn dense, incremental reward functions that are suitable for performing downstream reinforcement learning tasks, such as real-world robotic experiments. This is different from our approach, in that transfer learning is performed over the critic instead of the policy. Since pre-trained neural networks have displayed the potential to learn more sophisticated reward functions, it would be interesting to combine that reward with our approach through reward augmentation.

## 7. Conclusion

In this paper, we present a method to imitate complex behaviors and at the same time identify salient latent factors of variation in human decision-making. Discovering these latent factors does not require direct supervision beyond expert demonstrations, and the whole process can be trained end-to-end with standard policy optimization algorithms. We also introduce several techniques to successfully perform end-to-end imitation learning using visual inputs, including transfer learning and reward augmentation. Our experimental results in the TORCS simulator clearly show that our methods can automatically distinguish certain behaviors in human driving, while learning a policy that can imitate and even outperform the expert behavior, with visual information as the sole external input. We hope

that our work can further inspire end-to-end learning approaches to autonomous driving under more realistic scenarios.

Another compelling direction for future work is to explore how current mediated perception approaches could be combined with direct imitation learning through reward augmentation. Scene information could provide us with more powerful reinforcement signals, which are crucial to training machines that are better at driving than humans.

## Acknowledgements

Toyota Research Institute (TRI) provided funds to assist the authors with their research but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity.

## References

Abadi, Martín, Agarwal, Ashish, Barham, Paul, Brevdo, Eugene, Chen, Zhifeng, Citro, Craig, Corrado, Greg S, Davis, Andy, Dean, Jeffrey, Devin, Matthieu, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

Abbeel, Pieter and Ng, Andrew Y. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 1. ACM, 2004.

Aly, Mohamed. Real time detection of lane markers in urban streets. In *Intelligent Vehicles Symposium, 2008 IEEE*, pp. 7–12. IEEE, 2008.

Arjovsky, Martin, Chintala, Soumith, and Bottou, Léon. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.

Bloem, Michael and Bambos, Nicholas. Infinite time horizon maximum causal entropy inverse reinforcement learning. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pp. 4911–4916. IEEE, 2014.

Bojarski, Mariusz, Del Testa, Davide, Dworakowski, Daniel, Firner, Bernhard, Flepp, Beat, Goyal, Prasoon, Jackel, Lawrence D, Monfort, Mathew, Muller, Urs, Zhang, Jiakai, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.

Brockman, Greg, Cheung, Vicki, Pettersson, Ludwig, Schneider, Jonas, Schulman, John, Tang, Jie, and Zaremba, Wojciech. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

Chen, Chenyi, Seff, Ari, Kornhauser, Alain, and Xiao, Jianxiong. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2722–2730, 2015.

Chen, Xi, Duan, Yan, Houthooft, Rein, Schulman, John, Sutskever, Ilya, and Abbeel, Pieter. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2172–2180, 2016.

Geiger, Andreas, Lenz, Philip, Stiller, Christoph, and Urtasun, Raquel. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11): 1231–1237, 2013.

Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.

He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

Ho, Jonathan and Ermon, Stefano. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pp. 4565–4573, 2016.

Ho, Jonathan, Gupta, Jayesh K, and Ermon, Stefano. Model-free imitation learning with policy optimization. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.

Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

Kuefler, Alex, Morton, Jeremy, Wheeler, Tim, and Kochenderfer, Mykel. Imitating driver behavior with generative adversarial networks. *arXiv preprint arXiv:1701.06699*, 2017.

Lenz, Philip, Ziegler, Julius, Geiger, Andreas, and Roser, Martin. Sparse scene flow segmentation for moving object detection in urban environments. In *Intelligent Vehicles Symposium (IV), 2011 IEEE*, pp. 926–932. IEEE, 2011.

Levine, Sergey and Koltun, Vladlen. Guided policy search. In *ICML (3)*, pp. 1–9, 2013.

Lillicrap, Timothy P, Hunt, Jonathan J, Pritzel, Alexander, Heess, Nicolas, Erez, Tom, Tassa, Yuval, Silver, David, and Wierstra, Daan. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Pomerleau, Dean A. Alvinn, an autonomous land vehicle in a neural network. Technical report, Carnegie Mellon University, Computer Science Department, 1989.

Pomerleau, Dean A. Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 3(1):88–97, 1991.

Ratliff, Nathan D, Silver, David, and Bagnell, J Andrew. Learning to search: Functional gradient techniques for imitation learning. *Autonomous Robots*, 27(1):25–53, 2009.

Ross, Stéphane and Bagnell, Drew. Efficient reductions for imitation learning. In *AISTATS*, volume 3, pp. 3–5, 2010.

Ross, Stéphane, Gordon, Geoffrey J, and Bagnell, Drew. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, volume 1, pp. 6, 2011.

Russakovsky, Olga, Deng, Jia, Su, Hao, Krause, Jonathan, Satheesh, Sanjeev, Ma, Sean, Huang, Zhiheng, Karpathy, Andrej, Khosla, Aditya, Bernstein, Michael, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

Schaul, Tom, Quan, John, Antonoglou, Ioannis, and Silver, David. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*, 2015.

Schulman, John, Levine, Sergey, Abbeel, Pieter, Jordan, Michael I, and Moritz, Philipp. Trust region policy optimization. In *ICML*, pp. 1889–1897, 2015a.

Schulman, John, Moritz, Philipp, Levine, Sergey, Jordan, Michael, and Abbeel, Pieter. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015b.

Sermanet, Pierre, Xu, Kelvin, and Levine, Sergey. Unsupervised perceptual rewards for imitation learning. *arXiv preprint arXiv:1612.06699*, 2016.

Sharif Razavian, Ali, Azizpour, Hossein, Sullivan, Josephine, and Carlsson, Stefan. Cnn features off-the-shelf: an astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 806–813, 2014.

Silver, David, Huang, Aja, Maddison, Chris J, Guez, Arthur, Sifre, Laurent, Van Den Driessche, George, Schrittwieser, Julian, Antonoglou, Ioannis, Panneershelvam, Veda, Lanctot, Marc, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

Stadie, Bradly, Abbeel, Pieter, and Sutskever, Ilya. Third person imitation learning. In *ICLR*, 2017.

Syed, Umar, Bowling, Michael, and Schapire, Robert E. Apprenticeship learning using linear programming. In *Proceedings of the 25th international conference on Machine learning*, pp. 1032–1039. ACM, 2008.

Tamar, Aviv, Levine, Sergey, Abbeel, Pieter, WU, YI, and Thomas, Garrett. Value iteration networks. In *Advances in Neural Information Processing Systems*, pp. 2146–2154, 2016.

Tieleman, Tijmen and Hinton, Geoffrey. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2), 2012.

Williams, Ronald J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

Wymann, Bernhard, Espié, Eric, Guionneau, Christophe, Dimitrakakis, Christos, Coulom, Rémi, and Sumner, Andrew. Torcs, the open racing car simulator. *Software available at http://torcs. sourceforge. net*, 2000.

Yosinski, Jason, Clune, Jeff, Bengio, Yoshua, and Lipson, Hod. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, pp. 3320–3328, 2014.

Ziebart, Brian D, Maas, Andrew L, Bagnell, J Andrew, and Dey, Anind K. Maximum entropy inverse reinforcement learning. In *AAAI*, volume 8, pp. 1433–1438. Chicago, IL, USA, 2008.